arXiv:1611.07485v1 [cs.CV] 22 Nov 2016

Scene Labeling using Recurrent Neural Networks with Explicit Long Range Contextual Dependency

Qiangui Huang[1], Weiyue Wang[1], Kevin Zhou[2], Suya You[1], and Ulrich Neumann[1]

[1]University of Southern California, Los Angeles, CA
[2]Siemens Healthcare Technology Center, Princeton, NJ

Abstract

Spatial contextual dependencies are crucial for scene labeling problems. Recurrent neural network (RNN) is one of state-of-the-art methods for modeling contextual dependencies. However, RNNs are fundamentally designed for sequential data, not spatial data. This work shows that directly applying traditional RNN architectures, which unfold a 2D lattice grid into a sequence, is not sufficient to model structure dependencies in images due to the "impact vanishing" problem. A new RNN unit with Explicit Long-range Conditioning (RNN-ELC) is designed to overcome this problem. Based on this new RNN-ELC unit, a novel neural network architecture is built for scene labeling tasks. This architecture achieves state-of-the-art performances on several standard scene labeling datasets. Comprehensive experiments demonstrate that scene labeling tasks benefit a lot from the explicit long range contextual dependencies encoded in our algorithm.

1. Introduction

Scene labeling is a fundamental task in computer vision. Its goal is to assign one of many pre-defined category labels to each pixel in an image. It is usually modeled as a pixel-wise multi-class classification problem. Traditional methods rely on hand-crafted features like TextonBoost [37] and probabilistic graphical models (conditional random fileds [19], .etc.). Recently, many works have applied deep neural networks to this task and achieved significant improvements [7, 30, 50, 29].

Most of start-of-the-art scene labeling methods are based on convolutional neural networks (CNNs). CNNs are capable of learning scale-invariant discriminative features for images. These features have proven more powerful than traditional hand-crafted features on many computer vision tasks [20, 35, 30]. Specially designed CNN architectures [7, 30, 1, 49, 50] have shown superior performance on scene labeling by using end-to-end training.

However, CNNs have challenges in dealing with local textures and fine structures in images and tend to over-segment or under-segment objects in images. To accurately segment small components and detect object boundaries, long range contextual context in images can be helpful. Probabilistic graphical models such as conditional random fields (CRFs) [19] have been exploited to capture structural dependencies in images. Encouraging improvements have been achieved by jointly end-to-end training CNNs and CRFs [50, 29].

Recurrent neural networks (RNNs) have proven effective in modeling data dependencies in domains such as natural language processing and speech recognition [11, 33]. There are also some attempts at applying RNNs to images [4, 45, 39, 24]. Because the spatial relationship among pixels in 2D image data is fundamentally different from the temporal relationship in the 1D data in NLP or speech recognition, variants of RNN architectures [17, 39, 45, 43, 12, 25, 42] have been proposed to handle 2D image data, which typically involve unfolding a 2D lattice grid into a 1D sequence.. Nevertheless, these methods only explicitly condition on variables in a short range. As shown later, due to the "impact vanishing" problem, a short range conditioning does not effectively capture long range contextual dependencies among pixels that are actual spatial neighbors in images.

First, this paper describes the "impact vanishing" problem happening when traditional RNNs are applied to image data. Then, a new RNN unit with Explicit Long-range Conditioning (RNN-ELC) is developped to overcome this problem. Compared with existing works [17, 39, 45, 43, 12, 25, 42], RNN-ELC effectively captures long range contextual dependency in images. In the RNN-ELC unit, the present variable is explicitly conditioned on multiple sequentially distant but contextually related variables. Intuitively, it is adding skip connection between hidden states

belonging to distant variables. Adding skip connections in RNNs to help learn long term dependency first appears in [26]. Our method generalizes the idea to model more complex 2D spatial dependencies. The RNN-ELC unit is applicable to both raw pixels and image features. It can be naturally integrated in CNNs, thereby enabling joint end-to-end training.

Our main contributions include: 1) A descriptive study of the "impact vanishing" problem, which exists in traditional RNN units when they are applied to images; 2) A new RNN unit with Explicit Long-range Conditioning to alleviate the "impact vanishing" problem; 3) A novel scene labeling algorithm based on the new RNN-ELC unit; and 4) Improved performances on several standard scene labeling datasets.

2. Related Work

Scene labeling is one of the most challenging problems in computer vision as it involves several tasks in different levels. Recently, convolutional neural networks based methods achieved great success in this task. Farabet *et al.* [9] made one the earliest attempts at applying hierarchical features produced by CNNs to scene labeling. Eigen *et al.* [7]. designed a multi-scale convolutional architecture to jointly segment images, predict depth, and estimate normals for an image. Long *et al.* [30] applied fully convolutional network (FCN) for this task which can be obtained by simply surgerying powerful pre-trained deep CNNs and putting a deconvolution layer on top of it. Noh *et al.* [32] also used deconvolution layers for image segmentation. They adopted an encoder-decoder architecture, where encoder part consists of convolution and pooling operations and decoder part consists of deconvolution and unpooling operations. Badrinarayanan *et al.* [1] designed a similar architecture named SegNet. However, in SegNet, the decoder used stored pooling indexes of corresponding encoder to perform non-linear upsampling. In [49], Yu and Koltun developed a new dilated convolutional module to preserve multi-scale contextual information for image segmentation.

Although CNN based methods introduced powerful scale-invariant features for scene labeling, they performed poorly in preserving local textures and fine structures in predictions. These problems were addressed by combining CNNs with probabilistic graphical models such as conditional random fields (CRFs). Chen *et al.* [5] suggested to put a fully connected CRF [19] on top of FCN to capture structural dependencies in images. Their method achieved encouraging results but it treated CNN and CRF separately. Zheng *et al.* [50] showed that CNN and CRF can be jointly trained by passing the inference errors of CRFs back to CNN. Liu *et al.* [29] improved [50] by introducing a more complex pairwise term for CRF. These methods further advanced the scene labeling problem as CRFs can effectively utilize the unary features produced by CNN and capture pair-wise relations between pixels in images.

RNNs is another category of neural networks that have achieved tremendous success in many areas such as speech recognition [11] and natural language processing [33]. By implicitly conditioning on all previous variables, RNNs have proven effective in modeling internal dependencies in data. There of a rich literature of using RNNs for image related tasks [24, 47, 21, 34, 8].

Liang *et al.* [23] designed a graph LSTM to deal with image data. However, their method is built on superpixels over images, which is computationally expensive and is not directly applicable to image features. Byeon *et al.* [4] developed a 2D Long-Short Term Memory network for scene labeling. Their methods divided an input image into non-overlapping patches and then sequentially fed them into LSTM networks in four different orders. [44] built a RNN segmentation algorithm based on recently proposed ReNet [45]. Their idea was to to alternatively sweep an image in different directions and then sequentially input each row (or column) into a RNN. Shuai *et al.* [38, 39] designed a quad-directional 2D RNN architecture for scene labeling, where each pixel was connected to its 8 nearest neighbors.

Long range contextual dependencies are usually desired for scene labeling tasks. In previous works [17, 39, 45, 43, 12, 25, 42], the present variable only explicitly conditions on variables within a short range such as its 8 nearest neighbors. However, as shown later, short range conditioning does not capture long range structural dependencies in images. Our labeling algorithm is a generalized version of [44, 38, 39]. Compared with them, our model is built upon a new RNN unit with explicit long range contextual conditioning. Experimental results demonstrate our algorithm effectively captures structural dependencies in images and thus yields better performances in scene labeling problems.

3. Recurrent Neural Networks with Explicit Long-range Conditioning

3.1. Recurrent Neural Network

Recurrent neural networks (RNNs) are a rich class of neural network models that have been widely used to model contextual dependencies in sequential data. The vanilla RNN unit has two types of dense connections, namely, input-to-hidden and hidden-to-hidden connections. At each time step t, the output y^t conditions on the input at current time step x^t and the hidden state at previous time step h^{t-1}. Mathematically, given a sequence of input data $X = \{x^t, t = 1, ..., T\}$, the vanilla RNN unit models the probability of current time step output, $P(y^t|x^t, h^{t-1}, \theta)$, by the following equations:

$$h^t = \sigma_h(x^t W_x + h^{t-1} W_h + b_h), \quad (1)$$
$$y^t = \sigma_y(h^t W_y + b_y), \quad (2)$$

where $\theta = \{W_{x,h,y}, b_{h,y}\}$ is the parameter and σ_h and σ_y are the nonlinearity function of hidden and output layer, respectively. Though y^t explicitly conditions only on x^t and h^{t-1}, h^{t-1} explicitly conditions on previous input x^{t-1} and hidden state h^{t-2}; thus, y^t actually implicitly conditions on all previous inputs and hidden states. Intuitively, previous variables can influence their following variables by passing information through the hidden states.

However, vanilla RNNs have the notorious gradient vanishing or exploding problem when they are applied to learn long-term dependencies [2]. In practical applications, two types of *gated* RNNs are developed to avoid this problem, Long Short-Term Memory (LSTM) network [15] and Gated Recurrent Unit (GRU) network [6].

LSTM introduces self-loops to produce paths that help the gradient flow through a long sequence. It usually consists of four components to control information flow, an input gate, a forget gate, a cell gate, and an output gate. It uses the following equations to update its hidden states. Let i^t, f^t, c^t, and o^t denote the output of the input gate, the forget gate, the cell gate, and the output gate, respectively. $\odot$ denotes element-wise multiplication. $W_{xi,xf,xc,xo,hi,hf,hc,ho}$, $w_{ci,cf,co}$, and $b_{i,f,c,o}$ are parameters. $\sigma_{i,f,c,o}$ are nonlinearity functions.

$$i^t = \sigma_i(x^t W_{xi} + h^{t-1} W_{hi} + w_{ci} \odot c_{t-1} + b_i) \quad (3)$$
$$f^t = \sigma_f(x^t W_{xf} + h^{t-1} W_{hf} + w_{cf} \odot c_{t-1} + b_f) \quad (4)$$
$$c^t = f_t \odot c_{t-1} + i_t \odot \sigma_c(x^t W_{xc} + h_{t-1} W_{hc} + b_c) \quad (5)$$
$$o^t = \sigma_o(x^t W_{xo} + h_{t-1} W_{ho} + w_{co} \odot c_t + b_o) \quad (6)$$
$$h^t = o^t \odot \sigma_h(c^t) \quad (7)$$

Instead of using multiple gates in LSTM, GRUs use only one single gate to simultaneously control the forgetting factor and the decision to update the state unit. GRUs are usually composed of a reset gate and an update gate. They compute the output by the following equations with $W_{xr,hr,xu,hu,xc,hc}$ and $b_{r,u,c}$ being parameters.

$$r^t = \sigma_r(x^t W_{xr} + h^{t-1} W_{hr} + b_r) \quad (8)$$
$$u^t = \sigma_r(x^t W_{xu} + h^{t-1} W_{hu} + b_u) \quad (9)$$
$$c^t = \sigma_c(x^t W_{xc} + r^t \odot (h^{t-1} W_{hc}) + b_c) \quad (10)$$
$$h^t = (1 - u^t) \odot h^{t-1} + u^t \odot c^t \quad (11)$$

LSTMs and GRUs have been widely used in various applications as they effectively prevent gradients from vanishing or exploding [3]. However, as shown later, the length of dependency they can model is still limited for image related applications as the input sequence could be much longer than applications in other domains [11, 33].

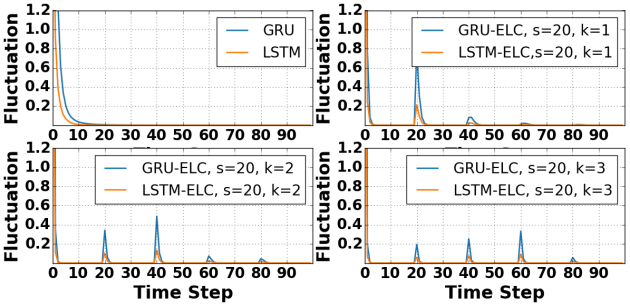

Figure 1. Illustration of "impact vanishing" problem in LSTMs and GRUs. Fluctuation denotes the difference caused by changing the first data x^1 in input sequence.

3.2. Impact Vanishing Problem

It is well known that long term dependency is hard to learn in RNNs. Beyond learning, here we show using a toy example that even with LSTM/GRU units, the length of long term dependency they can capture or model is still not large enough for image data.

Assume $X = \{x^t, t = 1, ..., T\}$ is an input sequence. $x^t \in R^{M \times N}$ is the input data at time step t and is generated from a continuous uniform distribution $U(0, 1)$. All weight parameters W in LSTMs and GRUs are initialized with a Guassian distribution $\mathcal{N}(0, 0.1)$ and bias parameters are set to zero by convention.

First, the entire sequence X is fed to LSTMs/GRUs that output $Y = \{y^t, t = 1, ..., T\}$. Then, the data at the first time step x^1 is replaced with $\hat{x}^1$ which is generated from same distribution $U(0, 1)$ and get a new input data sequence $\hat{X} = \{\hat{x}^1, x^2, ..., x^T\}$. Feed $\hat{X}$ to the same RNNs and get the new output $\hat{Y} = \{\hat{y}^1, \hat{y}^2, ..., \hat{y}^T\}$. It is reasonable to expect $\hat{y}^t$ should be different with y^t since different input at the first time step should result in different outputs in all following time steps. Our toy example calculates the following metric to measure the fluctuation at time step t caused by feeding different values of x^1:

$$F^t = \frac{1}{MN} \sum_{i=1,j=1}^{M,N} (y^t(i,j) - \hat{y}^t(i,j))^2 \quad (12)$$

We repeat this process by 20 times, collect all F^t, and report the mean of all 20 F^t in the top left plot in Fig.1. It shows that F^t drops dramatically in the first 10 time steps. When $t > 20$, F^t decreases to zero, which means that y^t stays unchanged when $t > 20$ regardless of the initial input x^1. Although h^t implicitly depends on h^1 and x^1, in

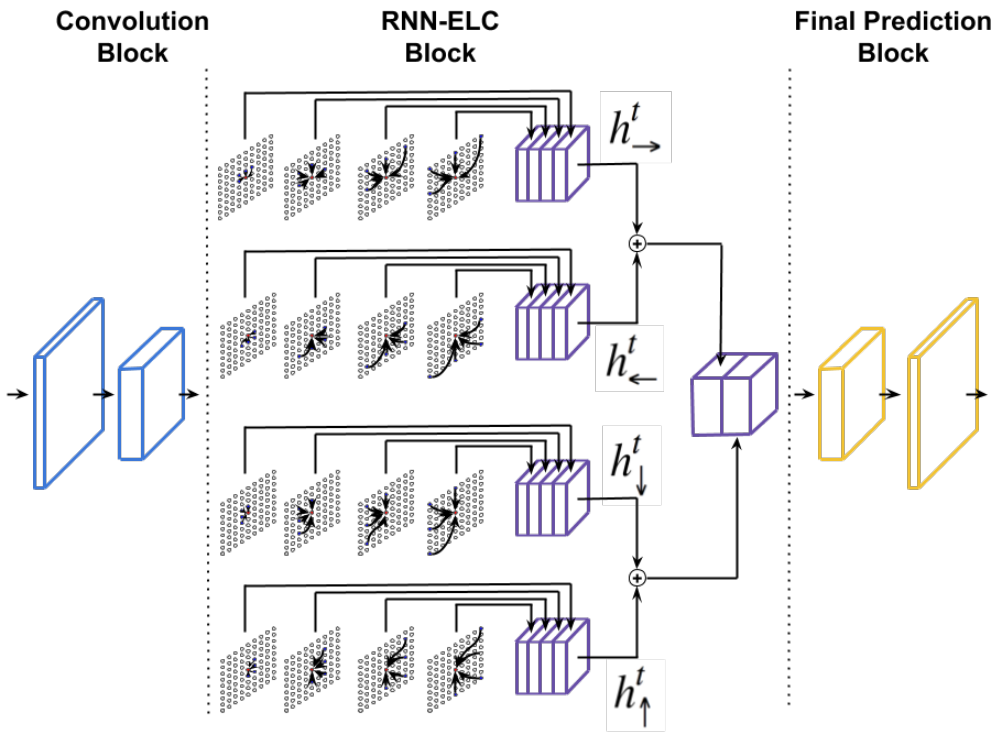

Figure 2. Illustration of our scene labeling algorithm described in Section 4 (with $k = 4$). Blue cuboids denote a convolution layer followed by a pooling layer. Purple cuboids denote concatenation operation. Yellow cuboids denote a unpooling layer followed by a convolution layer. In the RNN-ELC block, four parallel branches are used to calculate $h^t_{\rightarrow}$, $h^t_{\leftarrow}$, $h^t_{\downarrow}$, and $h^t_{\uparrow}$, respectively. Refer to equations (16)-(20) for more details.

practical, the dependency between variable 1 and variable t is broken when $t > 20$. This phenomenon is referred as "impact vanishing" problem in this paper. The explanation is straightforward and is similar to the reason for gradient vanishing problem [29]. According to the equations, when passing h^t to subsequent time steps, it needs to be multiplied by fixed weights over and over again. These weights usually fall in the range [-1 1], which makes the multiplication results rapidly vanish to zero and thus the variables at a sequantial distance from each other lose their dependency.

As RNNs are essentially designed for sequential data, image data needs to be sequentially fed into RNNs [17, 39, 45, 43, 12, 25, 42]. Some of the neighboring pixels could be positioned far away from each other (the distance is usually much larger than 20). This could break the contextual dependencies in images.

3.3. RNN units with Explicit Long-range Conditioning

The toy example in the previous section shows that implicitly conditioning on all previous variables is not powerful enough to model the long range dependency in RNNs. To tackle this problem, this work generalizes traditional RNN units to the RNN units with so-called Explicit Long-range Conditioning (RNN-ELC), in which sequentially distant but contextually related variables are explicitly conditioned on each other.

First assume that the variable at time step t is contextually related to both variable at $t-1$ and variable at $t-s$. We call s as a conditioning stride. Here is an example how this dependency assumption correlates to real world scenarios. Imagine there is an image with height h and width w and it is sequentially fed into an RNN from left to right and top to bottom. In this case, for the t^{th} pixel, besides the $(t-1)^{th}$ pixel, the one vertically above it is also high contextually related but is positioned at time step $(t-w)$. In this case, the variable t and variable $(t-w)$ are at a sequential distance from each other but contextually related from each other.

Take a vanilla RNN unit for example. The hidden state updating rule in equation (1) is changed into equation (13).

$$h^t = \sigma_h(x^t W_x + \frac{1}{2}(h^{t-1} + h^{t-s})W_h + b_h) \qquad (13)$$

$\frac{1}{2}$ is a constant term to keep h^t stay in the valid activation

area of σ_h. Now the RNN-ELC unit models the following conditional probability: $P(y^t|x^t, h^{t-1}, h^{t-s}, \theta)$. Note that no extra weights are introduced here.

It is straightforward to generalize equation (13) for an LSTM/GRU unit. This new conditioning rule is tested against the toy example using the same data and weights as in previous section by setting s as 20. F^t is plotted in the top right plot in Fig.1. Compared with traditional LSTM/GRU unit, F^t of the ELC-LSTM/GRU unit drops faster due to the constant term $\frac{1}{2}$. But, there is a strong peak around time step 20, which does not exist in tradition RNN units.

Intuitively, the peak means there is a strong dependency between the t^{th} and $(t-s)^{th}$ variable in RNN-ELC units. This is because the explicit conditioning on the $(t-s)^{th}$ variable has a huge impact on the output of the t^{th} variable.

In real world scenarios, a long range dependency is usually desired. The t^{th} pixel could be treated as related to all k pixels vertically above it. This means the t^{th} variable is contextually related to not only the $(t-s)^{th}$ variable but also $(t-2s)^{th}$,...,$(t-ks)^{th}$ variable. The top right plot in Fig.1 shows that a short range conditioning encoded by equation (13) is still not powerful enough to model this kind of long range dependency as the peak around time step 40 is relatively smaller and F^t decreases to near zero again when $t > 60$. So, equation (13) is further generalized to the equation (14)-(15) when a longer range dependency is desired. We call k as a conditioning scale.

$$h^t = \sigma_h(x^t W_x + H^t W_h + b_h) \tag{14}$$

$$H^t = \frac{1}{k+1}(h^{t-1} + \sum_{i=1}^{k} h^{t-i*s}) \tag{15}$$

Equations (14)-(15) are tested against the toy example again. s is set to be 20 and k is set to be 2 and 3. F^t is reported in the second row in Fig.1. It shows that by explicitly conditioning on related variables in a longer range, the length of valid dependency becomes longer. Next sections show how to exploit this property of RNN-ELC units and build a model for scene labeling which captures a desired long range contextual dependency.

4. RNN-ELC units for Scene Labeling

An overview of our scene labeling algorithm is presented in Fig.2. There are three blocks: convolution block, RNN-ELC block, and final prediction block. The convolution block encodes images into features. It is initialized by certain layers from VGG-16 network [41]. The final prediction block uses unpooling layers followed by convolution layers to create the feature maps of the same resolution as input images and outputs final predictions. This block is trained from scratch. The RNN-ELC block models contextual dependencies over features.

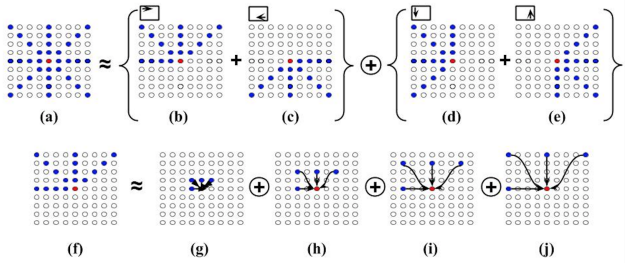

Figure 3. Illustration of dependency decomposition used in our scene labeling system. (a) denotes the spatial relation to be modeled. It is decomposed into a combination of (b)-(e) four parts. Arrow in each square box represents the sweep direction. In (b), image data is sweeped from left-to-right and top-to-bottom. (c) is from right-to-left and bottom-to-top. (d) is from top-to-bottom and left-to-right. (e) is from bottom-to-top and left-to-right. Second row presents how each of (b)-(e) is further decomposed. $\oplus$ implies concatenation operation here.

In this paper, the long range spatial relationship depicted in Fig.3(a) is considered in our RNN-ELC block. For the present variable (red dot in Fig.3(a)), its dependencies with k nearest variables (blue dots) in horizontal, vertical and diagonal directions are considered. However, when feeding 2D data into RNNs, it is infeasible to condition on all of them at the same time. Thus, the conditioning dependency is decomposed into four parts (Fig.3(b)-(e)). Each part feeds image data to RNNs in different orders (orders are specified by the arrows in the square box in Fig.3(b)-(e))). Let $h^t_{\rightarrow}$, $h^t_{\leftarrow}$, $h^t_{\downarrow}$, and $h^t_{\uparrow}$ denote the hidden states following the conditioning rules as in Fig.3(b), (c), (d), (e), respectively. Take $h^t_{\rightarrow}$ for example. Mathematically it adopts the following updating equations. Let w be the width of feature maps fed into the RNN-ELC block.

$$h^t_{\rightarrow} = \sigma_h(x^t W_x + H^t_{\rightarrow} W_h + b_h) \tag{16}$$

$$H^t_{\rightarrow} = \frac{1}{4k}\sum_{i=1}^{k}(h^{t-i}_{\rightarrow} + h^{t-iw}_{\rightarrow} + h^{t-(iw-i)}_{\rightarrow} + h^{t-(iw+i)}_{\rightarrow}) \tag{17}$$

Similarly we can get the updating equations for $h^t_{\leftarrow}$, $h^t_{\downarrow}$, and $h^t_{\uparrow}$. As $h^t_{\rightarrow}$ is complementary to $h^t_{\leftarrow}$ and $h^t_{\downarrow}$ is complementary to $h^t_{\uparrow}$, we sum them separately first, concatenate the sum results, and get the final results.

However, conditioning on too many terms in one unit limits the contribution of each of these terms due to the existence of the constant term $\frac{1}{4k}$. Thus, equations (16) - (17) are further decomposed into k separate parts as illustrated in the second row in Fig.3. Now, $h^t_{\rightarrow}$ is calculated by the following equations. $\oplus$ denotes the concatenation operation.

$$h^t_{\rightarrow} = h^t_{1\rightarrow} \oplus h^t_{2\rightarrow} \oplus ...h^t_{i\rightarrow} \oplus ...h^t_{k\rightarrow} \quad (18)$$

$$h^t_{i\rightarrow} = \sigma_h(x^t W_x + H^t_{i\rightarrow} W_h + b_h) \quad (19)$$

$$H^t_{i\rightarrow} = \frac{1}{4}(h^{t-i}_{i\rightarrow} + h^{t-iw}_{i\rightarrow} + h^{t-(iw-i)}_{i\rightarrow} + h^{t-(iw+i)}_{i\rightarrow}) \quad (20)$$

For summary, in our network, the convolution block encodes the input images into features first. The long range contextual dependencies are modeled in the RNN-ELC block according to equations (16) - (20). Finally, the final prediction block outputs prediction results. The key part of our algorithm is the RNN-ELC block. It is based on the newly designed RNN-ELC units which can effectively model desired long range contextual dependencies for images.

5. Experiment

In order to cover both outdoor and indoor scenes, we select three challenging datasets to test our method, namely the SiftFlow [27], NYUDv2 [40], and Stanford background [10] datasets. Two metrics are used for evaluation, namely global pixel accuracy (**Global**) and average per-class accuracy (**Class**). First, a full exploration of different architecture choices is conducted on the SiftFlow dataset. Then, state-of-the-art results are presented for the NYUDv2, and Stanford background dataset.

All experiments follow the same training protocol. Images with original size are used for training/testing. The initial learning rate is set as 0.001 and the poly learning policy is adopted to decrease the learning rate after every epoch. All models use recently developed Adam solver [18] for gradient descent training. The cross entropy loss is used as objective function. Median frequency balancing [7] is applied for comparison. Widely used data augmentation, cropping, flipping, and random jittering are adopted. The GRU-ELC unit is used in the RNN-ELC block for all experiments as it decays more slowly than the LSTM-ELC unit in our toy example (Fig.1).

5.1. The SiftFlow dataset

The SiftFlow dataset contains 2,688 images with 33 object categories. All images are of size 256×256 and captured in outdoor scenes like coast, highway, forest, city, etc. All experiments follow the same 2,488/200 training/testing split by convention.

First, our algorithm is built on top of the VGG-16 network. Convolutional layers up to *conv5_3* layer in VGG-16 are used as the convolution block in our network. Four pairs of unpooling and convolution layers are used in the final prediction block to generate the output predictions of the same resolution as input. Different settings are tested for the RNN-ELC block to demonstrate of the benefit of RNN-ELC units:

Method		Global (%)	Class (%)
Attent to rare class [48]		79.8	48.7
FCN-16s [30]		85.2	51.7
Eigen *et al.* [7]		86.8	46.4
Eigen *et al.* [7] *MB*		83.8	55.7
ParseNet [28]		86.8	52.0
RNN based	RCNN [22]	84.3	41.0
	DAG-RNN [39]	85.3	55.7
	Multi-Path [16]	86.9	56.5
	Attention [8]	86.9	57.7
conv3-LC-GRU-4		**87.8**	46.7
conv3-LC-GRU-4 *MB*		84.4	**64.0**

Table 2. Comparison with state-of-the-art on SiftFlow. Our method is compared against RNN related and other state-of-the-art methods. *MB* denotes this model is trained with median frequency balancing.

- **conv5-decoder**. The entire system is only composed of convolution block and final prediction block. No RNN units are used. It is similar to the encoder-decoder architectures in [32, 1].

- **conv5-GRU**. Only the traditional GRU unit is used in RNN-ELC block.

- **conv5-ELC-1**. GRU-ELC unit with $k = 1$ is used in RNN-ELC block. It is similar to the architecture presented in [39].

- **conv5-ELC-4**, GRU-ELC unit with $k = 4$ is used in RNN-ELC block.

Another two settings that are based on different convolution layers in the VGG-16 network are tested too:

- **conv4-ELC-4**. The system is built on top of convolutional layers up to the *conv4_3* layer in the VGG-16 network. The GRU-ELC unit with $k = 4$ is used in the RNN-ELC block. The final prediction block consists of three pairs of uppooling and convolution layers.

- **conv3-ELC-4**. The system is built on top of convolution layers up to the *conv3_3* layer in VGG-16 network. GRU-ELC unit with $k = 4$ is used in the RNN-ELC block. The final prediction block consists of two pairs of uppooling and convolution layers.

All the results are reported in Table 1, from which the following observations are made. First, experimental results show that without explicit long range dependencies, conv5-GRU performs much worse than conv5-decoder. After adding short range dependency (by setting $k = 1$), the

Method	No balancing		Median balancing	
	Global (%)	Class (%)	Global (%)	Class (%)
conv5-decoder	56.0	19.0	70.0	34.9
conv5-GRU	36.7	15.0	43.6	22.8
conv5-ELC-1	54.3	21.2	68.5	46.0
conv5-ELC-4	74.4	38.4	67.7	43.3
conv4-ELC-4	84.3	44.8	84.8	62.2
conv3-ELC-4	**87.8**	46.7	84.4	**64.0**

Table 1. Performance comparison of different choices for RNN-ELC block on SiftFlow

Input Information	Method	Global (%)	Class (%)
RGB	FCN 32s RGB [30]	60.0	42.2
RGB	conv4-ELC-4	**64.5**	41.4
RGB	conv4-ELC-4 *MB*	62.1	**45.5**
RGB+depth	Gupta *et al.*'13 [13]	59.1	28.4
RGB+depth	Gupta *et al.*'14 [14]	60.3	35.1
RGB+depth	FCN 32S RGBD [30]	61.5	42.4
RGB+depth	FCN 32S RGB-HHA [30]	64.3	44.9
RGB+depth	FCN 16S RGB-HHA [30]	65.4	46.1
RGB+depth+normal	Eigen *et al.* [7]	**65.6**	45.1
RGB+depth	Wang *et al.* [46]	-	**47.3**

Table 3. Comparison with state-of-the-art on NYU. *MB* denotes this model is trained with median frequency balancing.

Method	Global (%)	Class (%)
Sharma *et al.* [36]	81.9	73.6
Mostajabi *et al.* [31]	82.1	77.3
Liang *et al.* [22]	83.1	74.8
Multi-path [16]	86.6	79.6
conv3-ELC-4	**87.8**	81.1
conv3-ELC-4 *MB*	84.7	**81.5**

Table 4. Comparison with state-of-the-art on Stanford Background Dataset.*MB* denotes this model is trained with median frequency balancing.

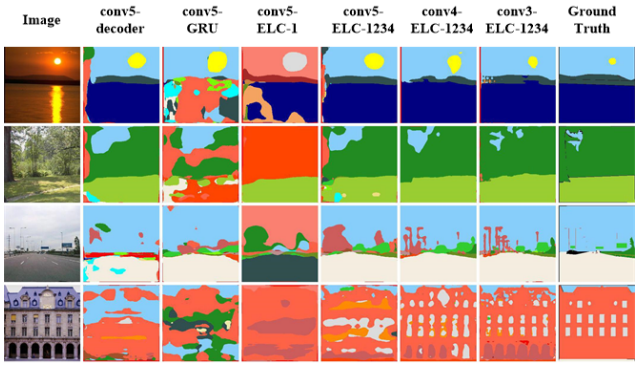

Figure 4. Comparison between different settings of RNN-ELC block in our system. All models are trained with median frequency balancing. Globally, conv5-ELC-1234, conv4-ELC-1234, and conv3-ELC-1234 all make reasonable prediction results. But conv3-ELC-1234 performs better in preserving local details such as the sun in first row, the sky in second row, the poles in the third row, and windows in the fourth row even some of them are not annotated in grounth truth.

performance is improved to a comparable level to conv5-decoder. Adding longer range dependency ($k = 4$) dramatically improves the labeling performances. This demonstrates that scene labeling tasks can benefit a lot from the long range contextual dependencies encoded in our RNN-ELC units. Note that using lower layers in the VGG-16 network can further improve the performance of our system. This is perhaps because the features from lower layers contain richer contextual details. Also median frequency balancing helps the model achieve better average per-class accuracy by sacrificing global pixel accuracy. Visual comparisons of the results produced by different RNN-ELC blocks are presented in Fig.4.

Our methods are also compared with other state-of-the-art results in Table 2. Comprehensive comparisons show that our methods can outperform both RNN based and other state-of-the-art methods. Especially the average class accuracy has been improved by near 6.3%! These results show that explicit long range contextual dependencies in RNN units help a lot for the scene labeling task. Results produced by our algorithm and FCN-8s [30] are visually compared in

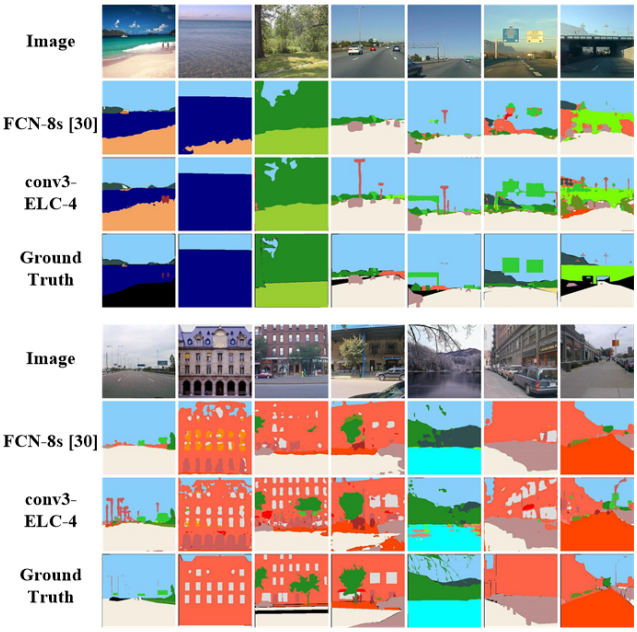

Figure 5. Comparisons of results produced by FCN-8s [30] and conv3-ELC-4 trained with median frequency balancing. conv3-ELC-4 can accurately predict small objects in scenes such the persons, poles, boards, and windows. Note that the black color denotes unknown categories.

Fig.5.

5.2. The NYUDv2 dataset

NYUDv2 is an RGB-D dataset containing 1,449 RGB and depth image pairs for indoor scenes. A standard split contains 795 training images and 654 testing images. 40 categories [13] have been widely used to test the performance of labeling algorithms. Besides RGB images, depth [30, 14] and normal information [7] can also be used for scene labeling. However only RGB images are used in our algorithm in order to get a clear sense about the capacity of our algorithm. The conv4-ELC-4 setting is applied for this dataset. The comparison between our method and other state-of-the-art methods are reported in Table 3. Although only RGB information is used in our method, our results are quite competitive against other methods that additionally use depth or normal information. Some sample results produced by our algorithm are compared with FCN-32s-RGB [30] and Eigen *et al.* [7] in Fig.6.

5.3. The Stanford background dataset

The Stanford background dataset is composed of 715 images with 8 object categories. The images are captured in outdoor scenes and most of them are of size 320×240. Following the standard protocol [22], a 5-fold cross validation is used for measuring the performance, each of them ran-

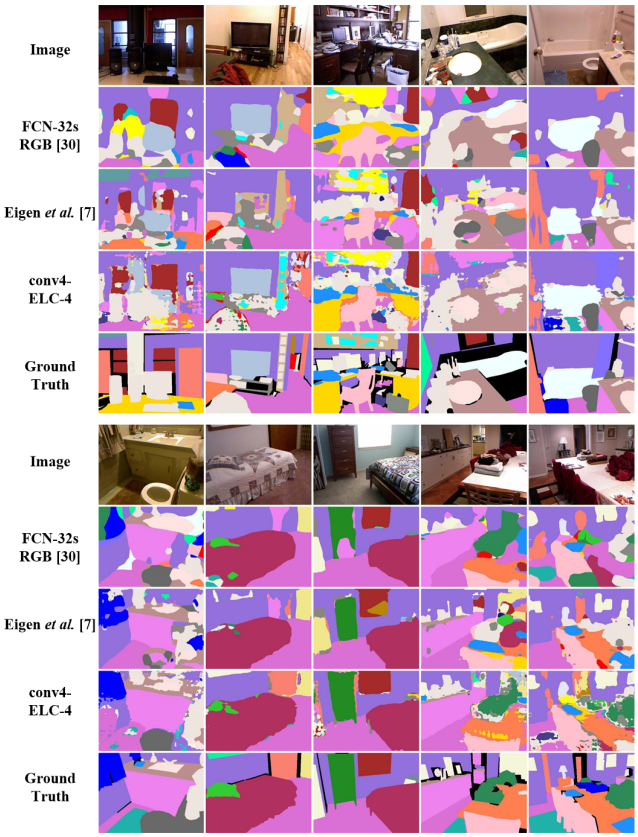

Figure 6. Comparisons of results produced by FCN-8s [30], Eigen *et al.* [7], and conv4-ELC-4 trained with median frequency balancing. Note that the black color denotes unknown categories.

domly selecting 572 images for training and 143 images for testing. The conv3-ELC-4 setting is adopted for this dataset. Results of our method are reported and compared with other state-of-the-art methods in Table 4. Our state-of-the-arts results demonstrate the effective of long range dependency for scene labeling problems.

6. Conclusion

Long range contextual dependencies are important for scene labeling tasks. RNNs are a class of neural network models that have proven effective in modeling internal data dependency in many areas. However, traditional RNN units that unfold a 2D lattice grid into a 1D sequence are not powerful enough to model the contextual dependencies in images due to the "impact vanishing" problem. A new RNN unit with Explicit Long-range Conditioning is designed to avoid this problem. Based on the RNN-ELC units, a novel scene labeling algorithm is developed in this paper. Various experimental results and comparisons with other state-of-the-art methods demonstrate that our algorithm can effectively capture long range contextual dependencies in images

and thus give better performances in scene labeling.

Potential directions for future work include: 1) Extending our scene labeling algorithm to take multi-modal input information (like depth or normal information); and 2) Applying our new RNN-ELC unit to other image related applications such as image inpainting and image generation.

References

[1] V. Badrinarayanan, A. Handa, and R. Cipolla. Segnet: A deep convolutional encoder-decoder architecture for robust semantic pixel-wise labelling. *arXiv preprint arXiv:1505.07293*, 2015.

[2] Y. Bengio, I. J. Goodfellow, and A. Courville. Deep learning. *An MIT Press book in preparation. Draft chapters available at http://www. iro. umontreal. ca/ bengioy/dlbook*, 2015.

[3] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994.

[4] W. Byeon, T. M. Breuel, F. Raue, and M. Liwicki. Scene labeling with lstm recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3547–3555, 2015.

[5] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014.

[6] K. Cho, B. Van Merriënboer, D. Bahdanau, and Y. Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.

[7] D. Eigen and R. Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2650–2658, 2015.

[8] H. Fan, X. Mei, D. Prokhorov, and H. Ling. Multi-level contextual rnns with attention model for scene labeling. *arXiv preprint arXiv:1607.02537*, 2016.

[9] C. Farabet, C. Couprie, L. Najman, and Y. LeCun. Learning hierarchical features for scene labeling. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1915–1929, 2013.

[10] S. Gould, R. Fulton, and D. Koller. Decomposing a scene into geometric and semantically consistent regions. In *2009 IEEE 12th international conference on computer vision*, pages 1–8. IEEE, 2009.

[11] A. Graves, A.-r. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 6645–6649. IEEE, 2013.

[12] A. Graves and J. Schmidhuber. Offline handwriting recognition with multidimensional recurrent neural networks. In *Advances in neural information processing systems*, pages 545–552, 2009.

[13] S. Gupta, P. Arbelaez, and J. Malik. Perceptual organization and recognition of indoor scenes from rgb-d images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 564–571, 2013.

[14] S. Gupta, R. Girshick, P. Arbeláez, and J. Malik. Learning rich features from rgb-d images for object detection and segmentation. In *European Conference on Computer Vision*, pages 345–360. Springer, 2014.

[15] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[16] X. Jin, Y. Chen, J. Feng, Z. Jie, and S. Yan. Multi-path feedback recurrent neural network for scene parsing. *arXiv preprint arXiv:1608.07706*, 2016.

[17] N. Kalchbrenner, I. Danihelka, and A. Graves. Grid long short-term memory. *arXiv preprint arXiv:1507.01526*, 2015.

[18] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[19] V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. *Adv. Neural Inf. Process. Syst*, 2011.

[20] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[21] Z. Li, Y. Gan, X. Liang, Y. Yu, H. Cheng, and L. Lin. Lstm-cf: Unifying context modeling and fusion with lstms for rgb-d scene labeling. In *European Conference on Computer Vision*, pages 541–557. Springer, 2016.

[22] M. Liang, X. Hu, and B. Zhang. Convolutional neural networks with intra-layer recurrent connections for scene labeling. In *Advances in Neural Information Processing Systems*, pages 937–945, 2015.

[23] X. Liang, X. Shen, J. Feng, L. Lin, and S. Yan. Semantic object parsing with graph lstm. *arXiv preprint arXiv:1603.07063*, 2016.

[24] X. Liang, X. Shen, D. Xiang, J. Feng, L. Lin, and S. Yan. Semantic object parsing with local-global long short-term memory. *arXiv preprint arXiv:1511.04510*, 2015.

[25] X. Liang, X. Shen, D. Xiang, J. Feng, L. Lin, and S. Yan. Semantic object parsing with local-global long short-term memory. *arXiv preprint arXiv:1511.04510*, 2015.

[26] T. Lin, B. G. Horne, P. Tino, and C. L. Giles. Learning long-term dependencies is not as difficult with narx recurrent neural networks. 1998.

[27] C. Liu, J. Yuen, and A. Torralba. Nonparametric scene parsing: Label transfer via dense scene alignment. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1972–1979. IEEE, 2009.

[28] W. Liu, A. Rabinovich, and A. C. Berg. Parsenet: Looking wider to see better. *arXiv preprint arXiv:1506.04579*, 2015.

[29] Z. Liu, X. Li, P. Luo, C.-C. Loy, and X. Tang. Semantic image segmentation via deep parsing network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1377–1385, 2015.

[30] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.

[31] M. Mostajabi, P. Yadollahpour, and G. Shakhnarovich. Feedforward semantic segmentation with zoom-out features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3376–3385, 2015.

[32] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1520–1528, 2015.

[33] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. *ICML (3)*, 28:1310–1318, 2013.

[34] P. H. Pinheiro and R. Collobert. Recurrent convolutional neural networks for scene labeling. In *ICML*, pages 82–90, 2014.

[35] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.

[36] A. Sharma, O. Tuzel, and M.-Y. Liu. Recursive context propagation network for semantic scene labeling. In *Advances in Neural Information Processing Systems*, pages 2447–2455, 2014.

[37] J. Shotton, J. Winn, C. Rother, and A. Criminisi. Textonboost: Joint appearance, shape and context modeling for multi-class object recognition and segmentation. In *European conference on computer vision*, pages 1–15. Springer, 2006.

[38] B. Shuai, Z. Zuo, and G. Wang. Quaddirectional 2d-recurrent neural networks for image labeling. *IEEE Signal Processing Letters*, 22(11):1990–1994, 2015.

[39] B. Shuai, Z. Zuo, G. Wang, and B. Wang. Dag-recurrent neural networks for scene labeling. *arXiv preprint arXiv:1509.00552*, 2015.

[40] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from rgbd images. In *European Conference on Computer Vision*, pages 746–760. Springer, 2012.

[41] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[42] L. Theis and M. Bethge. Generative image modeling using spatial lstms. In *Advances in Neural Information Processing Systems*, pages 1927–1935, 2015.

[43] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016.

[44] F. Visin, M. Ciccone, A. Romero, K. Kastner, K. Cho, Y. Bengio, M. Matteucci, and A. Courville. Reseg: A recurrent neural network-based model for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 41–48, 2016.

[45] F. Visin, K. Kastner, K. Cho, M. Matteucci, A. Courville, and Y. Bengio. Renet: A recurrent neural network based alternative to convolutional networks. *arXiv preprint arXiv:1505.00393*, 2015.

[46] J. Wang, Z. Wang, D. Tao, S. See, and G. Wang. Learning common and specific features for rgb-d semantic segmentation with deconvolutional networks. In *European Conference on Computer Vision*, pages 664–679. Springer, 2016.

[47] Z. Yan, H. Zhang, Y. Jia, T. Breuel, and Y. Yu. Combining the best of convolutional layers and recurrent layers: A hybrid network for semantic segmentation. *arXiv preprint arXiv:1603.04871*, 2016.

[48] J. Yang, B. Price, S. Cohen, and M.-H. Yang. Context driven scene parsing with attention to rare classes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3294–3301, 2014.

[49] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.

[50] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1529–1537, 2015.